

Towards Better IncomLDL: We Are Unaware of Hidden Labels in Advance

Jiecheng Jiang^{1,2*}, Jiawei Tang^{1*}, Jiahao Jiang¹, Hui Liu³, Junhui Hou^{4†}, Yuheng Jia^{1,3,5†}

¹School of Computer Science and Engineering, Southeast University, Nanjing 210096, China

²School of Computing, National University of Singapore, Singapore

³School of Computing Information Sciences, Saint Francis University, Hong Kong, China

⁴Department of Computer Science, City University of Hong Kong, Hong Kong, China

⁵Key Laboratory of New Generation Artificial Intelligence Technology and Its

Interdisciplinary Applications (Southeast University), Ministry of Education, China

{jcjiang, jwtang, jhjiang, yhjia}@seu.edu.cn, hliu99-c@my.cityu.edu.hk, jh.hou@cityu.edu.hk

Abstract

Label distribution learning (LDL) is a novel paradigm that describe the samples by label distribution of a sample. However, acquiring LDL dataset is costly and time-consuming, which leads to the birth of incomplete label distribution learning (IncomLDL). All the previous IncomLDL methods set the description degrees of “missing” labels in an instance to 0, but remains those of other labels unchanged. This setting is unrealistic because when certain labels are missing, the degrees of the remaining labels will increase accordingly. We fix this unrealistic setting in IncomLDL and raise a new problem: LDL with hidden labels (HidLDL), which aims to recover a complete label distribution from a real-world incomplete label distribution where certain labels in an instance are omitted during annotation. To solve this challenging problem, we discover the significance of proportional information of the observed labels and capture it by an innovative constraint to utilize it during the optimization process. We simultaneously use local feature similarity and the global low-rank structure to reveal the mysterious veil of hidden labels. Moreover, we **theoretically** give the recovery bound of our method, proving the feasibility of our method in learning from hidden labels. Extensive recovery and predictive experiments on various datasets prove the superiority of our method to state-of-the-art LDL and IncomLDL methods.

Code — <https://github.com/Trisitana/HidLDL>

1 Introduction

Learning with ambiguity has gained notable attention in machine learning community recently (Geng 2016). In classical machine learning, single-label learning refers to the process of assigning exactly one label to an instance. Multi-label learning (Tsoumakas and Katakis 2008; Zhang and Zhou 2014; Liu et al. 2021), on the other hand, assigns multiple labels to a single instance. However, in practical applications, for the same instance, we need to consider the relative importance of their labels. Label distribution learning (LDL) (Geng 2016) is a novel machine learning paradigm

to solve the above issue by modeling the importances of multiple labels for each instance. In recent years, LDL have been widely used in many applications, such as age estimation (Geng, Yin, and Zhou 2013; Gao et al. 2018), emotion recognition (Jia et al. 2019b; Chen et al. 2020), facial beauty (Liang et al. 2018), crowd opinion prediction (Geng and Hou 2015), medical diagnosis (Wang et al. 2022), and scene text detection (Ren and Geng 2017; Ma et al. 2023) etc.

In LDL, each instance x is annotated by a real value vector d_x , i.e., label distribution (LD). Each component of the LD vector d_x^y represents the description degree of the label y to the instance x . Assume $d_x^y \in [0, 1]$ and the label set is complete, which means that all the labels in the set can always fully describe the instance. Then, we have $\sum_y d_x^y = 1$.

Despite LDL being successfully applied to many scenarios in recent years, existing LDL methods are incapable of handling incomplete supervised information. In practice, the description degrees are often provided by human annotators, which can be extremely costly in terms of labor and time when dealing with a large dataset. Meanwhile, data corruption and incomplete annotations may also occur. To this end, incomplete label distribution learning (IncomLDL) (Xu and Zhou 2017) was proposed to address the problem of incomplete supervised data. In the previous research of IncomLDL, a predefined set of labels is given. The description degrees of some labels for a given instance are randomly set to 0, and those of other labels still remain unchanged.

However, the setting of IncomLDL is **unreasonable** and **unrealistic**. In real situations, when certain labels are missing (their description degrees set to 0 in IncomLDL), degrees of the remaining labels should increase accordingly instead of remaining unchanged. Below, we will use an illustration to explain this unreasonable setting inherited in IncomLDL.

Fig. 1a shows a scene image containing 6 elements, and its LD $d_x = \{0.52, 0.14, 0.07, 0.15, 0.03, 0.18\}$ is presented in Fig. 1b. In IncomLDL, some of the description degrees are randomly set to 0 so that incomplete data of LD are formed into $d_x = \{0.52, 0, 0.07, 0.15, 0.03, 0\}$ (Fig. 1c). This is an approximation of real-world missing data. However, unless a portion of the data in the database is lost, this setting has a significant distance from the real world. Suppose we are the annotators. When annotating the image in Fig. 1a, we

*These authors contributed equally.

†Corresponding authors.




Figure 1: (a) A scene image containing 6 elements. (b) The label distribution (LD) of the image. (c) The observed LD in IncomLDL, with gray indicating missing (unobserved) labels. The sum of the description degrees of the observed labels is not 1. (d) The observed LD in HidLDL. For the unobserved labels in (c), the corresponding description degrees are 0, and the sum of the description degrees of the observed labels is 1. HidLDL is more intuitive and realistic, as the observed labels should occupy all description degree.

accidentally omit the labels for “cloud” and “lawn” (such omitted labels are referred to hereafter as hidden labels). In this case, the remaining labels ought to form a new comprehensive description of the instance, and the sum of their description degrees should equal 1. As illustrated in Fig. 1d, with “cloud” and “lawn” hidden from the annotator’s view, the description degrees of the other labels will increase to compensate and still manage to provide a complete description. Therefore, in IncomLDL, it is not a viable practice to simply set description degrees of hidden labels to 0 while keeping the description of the observed labels unchanged.

Based on the aforementioned illustration, we contend that current research on IncomLDL does not adequately reflect the true nature of incompleteness and leaves considerable room for refinement. We propose a new setting called LDL with hidden labels (HidLDL), which is closer to real situations. **We then propose a novel method that pays close attention to the crucial proportion information of observed labels, which is our most outstanding innovation.** By exploiting local feature similarity and the global low-rank structure to capture both local and global dependencies, we strive for further improvement. Moreover, we **theoretically** give the recovery bound of our method, proving the feasibility of our method in learning from hidden labels. Experimental results have shown that our method has significantly better performance compared to existing methods. The main contributions can be summarized as follows.

- We notice the unreasonable aspect of all previous incomplete LDL research, and propose a novel setting, abbreviated as HidLDL, which is closer to real situations.
- We discover the need to fully utilize the proportion information of observed labels in the new HidLDL problem. During optimization, apart from simultaneously using both local feature similarity and global low-rank structure, we innovatively use constraints to ensure the correct use of proportion information.
- We theoretically give the recovery bound of our method, in which the optimization objective and constraints of our

method help determine the recovery bound, proving the feasibility of our method in learning from hidden labels.

- We empirically verify the effectiveness of our method on 12 real-world datasets in both recovery and predictive tasks. Experimental results show that our method outperforms state-of-the-art LDL and IncomLDL methods, with the improvements being statistically significant in numerous cases.

2 Related Work

2.1 Label Distribution Learning

Label distribution learning (LDL) (Xie et al. 2026; Tang and Jia 2025; Jia, Tang, and Jiang 2024; Kou et al. 2025b,a) was first proposed to solve the facial age estimation problem. Later on, (Geng 2016) discovered that in some real applications, the distribution across all labels is more desirable than the association of a single label to an instance. Since LDL was introduced, numerous methods have been proposed. These methods can be roughly divided into three categories: problem transformation, algorithm adaptation and specialized algorithms. Yang (Yang et al. 2015) proposed a deep learning based LDL method. Two classic methods SA-IIS and BFGS-LLD were then proposed in (Geng 2016). LDLLC (Jia et al. 2018) maps label distances to parameter matrix column distances and uses Pearson correlation to capture global label relationships. GD-LDL-SCL and Adam-LDL-SCL (Jia et al. 2019a) were designed to exploit local correlation in different groups which helps solve LDL problems. LDL-LDM (Wang and Geng 2021) was proposed to exploit label correlation manifold which is helpful to alleviate the overwhelm of LDL outspace. In (Wen et al. 2023), the authors proposed an auxiliary MLL process integrated into the LDL framework, which captures low-rank label correlations within this additional MLL component. Also, a new paradigm based on LDL, called OLDL (Wen et al. 2023), was proposed focusing on the internal sequential patterns of labels.

2.2 Incomplete Label Distribution Learning

The above mentioned LDL methods all assume that in the training set, each sample is annotated with a full label distribution. But in reality, the labels of some samples are either not observed, or it is expensive to annotate complete training samples (Li, Li, and Jia 2025; Jia et al. 2024, 2022). To this end, incomplete label distribution learning (IncomLDL) was introduced by (Xu and Zhou 2017) to address annotation incompleteness. To solve the LDL problem when given incomplete supervised information, they suggested a low-rank model to utilize the correlation between labels and proposed IncomLDL-prox and IncomLDL-admm. Adapting existing LDL methods to the scenario of missing annotations is another intuitive strategy for addressing incomplete LDL tasks. In (Li and Chen 2024), the authors designed a weighting scheme of label degrees to utilize prior information of label distribution itself. In (Wang and Geng 2021), the authors contended that the assumption of low-rank might not always be valid. They proposed an alternative model where the predictions are constrained to lie on a shared manifold, with the manifold’s structure capturing the correlations between labels. Recently, in (Xu et al. 2025), the authors proposed an Incomplete Label Distribution Learning method via correlation decomposition by assuming a globally low-rank structure and a locally sparse structure. They developed an alternating solution with the accelerated proximal gradient descent method for optimization (IncomLDL-LCD).

3 LDL with Hidden Labels

3.1 Problem Setting

Let $\mathbf{x} \in R^d$ denote a feature vector, $\mathbf{X} = [\mathbf{x}_1; \mathbf{x}_2; \dots; \mathbf{x}_n] \in R^{n \times d}$ denote the feature matrix and $\mathcal{Y} = \{y_1, y_2, \dots, y_m\}$ denote the complete set of labels, where d , n and m represent the dimension of features, the number of samples, and the number of labels. In LDL, the description degree of the label y to the instance \mathbf{x} is denoted by $d_{\mathbf{x}}^y$, which satisfies $d_{\mathbf{x}}^y \in [0, 1]$ and $\sum_y d_{\mathbf{x}}^y = 1$, the label distribution of \mathbf{x}_i is denoted by $\mathbf{d}_i^g = (d_{\mathbf{x}_i}^{y_1}, d_{\mathbf{x}_i}^{y_2}, \dots, d_{\mathbf{x}_i}^{y_m}) \in R^m$, and the ground-truth label distribution matrix is denoted by $\mathbf{D}^g = [\mathbf{d}_1^g; \mathbf{d}_2^g; \dots; \mathbf{d}_n^g] \in R^{n \times m}$.

In IncomLDL (Xu and Zhou 2017), for each sample \mathbf{x} , there exists a possibility that some entries in $d_{\mathbf{x}}^y$ may not be observed, aiming to simulate the real-world scenarios. Specially, let $\mathbf{M} \in R^{n \times m}$ denote the indices of the observed labels in \mathbf{D}^g :

$$M_{ij} = \begin{cases} 1 & \text{if } D_{ij}^g \text{ is observed,} \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Then, the observed label distribution matrix in IncomLDL is defined as,

$$D_{ij}^o = \begin{cases} D_{ij}^g & \text{if } M_{ij} = 1, \\ 0 & \text{if } M_{ij} = 0. \end{cases} \quad (2)$$

However, this approach is not reasonable and has been discussed in detail in the previous context.

Therefore, we propose a new setting called HidLDL. Specifically, in HidLDL, due to some labels not being observed at the beginning, in the eyes of the annotator, other

labels will occupy the degree of these hidden labels. We assume that the degrees of the observed labels in \mathbf{D}^g will amplify proportionally. Based on this instinctive assumption, a more reasonable \mathbf{D}^o is defined as,

$$D_{ij}^o = \begin{cases} \frac{D_{ij}^g}{\sum_{k=1}^m D_{ik}^g \cdot M_{ik}} & \text{if } M_{ij} = 1, \\ 0 & \text{if } M_{ij} = 0. \end{cases} \quad (3)$$

The goal of HidLDL is to recover a label distribution matrix $\mathbf{D} = [\mathbf{d}_1; \mathbf{d}_2; \dots; \mathbf{d}_n] \in R^{n \times m}$, which is as close as possible to the ground-truth label distribution. Due to the inaccuracy of the observed label distribution, HidLDL receives more ambiguity compared to IncomLDL and requires the use of new approach to solve it.

3.2 Proposed Method

In order to recover hidden labels, a natural idea is to learn from other instances containing these hidden labels. Based on the smoothness assumption (Zhu 2005), in the feature space, instances close to each other tend to share a common label. Even if the observed degree of labels is not accurate in HidLDL, the nearby points can still learn from each other.

Graph Dependency To capture graph dependency, we use K-nearest neighbors algorithm to determine whether \mathbf{x}_i and \mathbf{x}_j are connected. Setting the neighbor number to the label number m , \mathbf{x}_i and \mathbf{x}_j are connected if \mathbf{x}_i is among K-nearest neighbors of \mathbf{x}_j or vice versa. At first, the similarity matrix \mathbf{A} is set to $\mathbf{0}_{n \times n}$, indicating that there are no initial dependencies. Then, if \mathbf{x}_i and \mathbf{x}_j are connected, A_{ij} is updated in Eq. (4):

$$A_{ij} = \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{2\sigma^2}\right), \quad (4)$$

where σ is the bandwidth parameter of the Gaussian kernel. If \mathbf{x}_i and \mathbf{x}_j are closely connected, then \mathbf{d}_i and \mathbf{d}_j should be close to each other. This idea leads to the following optimization objectives $\Omega(\mathbf{D})$:

$$\Omega(\mathbf{D}) = \sum_{i,j} A_{ij} \|\mathbf{d}_i - \mathbf{d}_j\|_2^2 = \text{tr}(\mathbf{D}^\top \mathbf{G} \mathbf{D}), \quad (5)$$

where $\mathbf{G} = \hat{\mathbf{A}} - \mathbf{A}$ is the graph Laplacian and $\hat{\mathbf{A}}$ is a diagonal matrix whose elements are $\hat{A}_{ii} = \sum_{j=1}^n A_{ij}$.

Low-rank Assumption Global low-rank assumption has good performance in handling incomplete supervision information (Xu and Zhou 2017). To further mine the correlation between labels, we assume that \mathbf{D} is low-rank, i.e., we hope that the trace norm of the recovered matrix \mathbf{D} is as small as possible. Therefore, the trace norm of \mathbf{D} (i.e., $\|\mathbf{D}\|_*$) is added as an optimization term, which is the sum of all singular values of the matrix.

Proposed Proportional Constraint The most important thing in HidLDL is how to utilize the observed matrix \mathbf{D}^o . In (Li and Chen 2024), Frobenius norm is used to minimize the difference in those observed positions, however, in HidLDL, this may not work because observed degrees of labels no longer directly reflect their final description of the instance. Specifically, we must only use the proportional relationship between the degrees of the observed labels to avoid introducing more noise. Even if we try to incorporate this proportion information of \mathbf{D} into the final optimization objective, it not only introduces additional hyperparameters but also brings in more non-smooth terms. Thus, we define the following constraint set:

$$\text{Cons} = \{ \mathbf{Z} \in R^{n \times m} \mid \forall i \in \{1, 2, \dots, n\}, \exists k_i \in R \text{ s.t. } k_i \cdot \mathbf{d}_i^o = \mathbf{z}_i \odot \mathbf{m}_i \}, \quad (6)$$

where the i -th row of \mathbf{Z} is denoted as \mathbf{z}_i , so is \mathbf{m}_i and \mathbf{d}_i^o . We constrain $\mathbf{D} \in \text{Cons}$. We expect that each row of the recovered label matrix, after masking, is proportional to the initial observed matrix, with a scaling coefficient of k_i . This setting makes sense even when \mathbf{D}_{ij}^o is 0, thus increasing the generality of our model.

Note that the recovered label distribution needs to be in the probability simplex. We add constraints $\mathbf{D} \times \mathbf{1}_m = \mathbf{1}_n$ and $\mathbf{D} \geq \mathbf{0}_{n \times m}$ to guarantee that all entries are negative and the sum of each row's entries is 1. Considering all the factors mentioned above, the final objective function of our method can be written as follows:

$$\min_{\mathbf{D}} \frac{1}{2} \text{tr}(\mathbf{D}^\top \mathbf{G} \mathbf{D}) + \alpha \|\mathbf{D}\|_* \quad (7)$$

s.t. $\mathbf{D} \times \mathbf{1}_m = \mathbf{1}_n, \mathbf{D} \geq \mathbf{0}_{n \times m}, \mathbf{D} \in \text{Cons},$

where $\text{tr}(\cdot)$ denotes the trace of a matrix, $\|\cdot\|_*$ denotes the trace norm of the matrix, and α serves as a regularization parameter that controls the trade-off between graph dependency and the trace norm.

3.3 Theoretical Analysis

In this section, we provide the recovery bound of our method.

Lemma 1. *For i -th instance, the recovery bound of it can be represented by the error between the ground-truth scaling coefficient $k_i^g = \frac{d_{ij}^o}{d_{ij}^g}$ and the recovered scaling coefficient k_i . So, the target is to prove that*

$$(k_i^g - k_i)^2 \leq \epsilon_i, \quad (8)$$

where ϵ_i is a small number.

Theorem 1. *Under certain assumptions, the error of the scaling coefficients can be expressed as*

$$k_i^g - k_i \leq \frac{1 - \sum_{k=1}^m D_{ik} M_{ik}}{\sum_{k=1}^m D_{ik}^g M_{ik}}. \quad (9)$$

Let $\sigma_i = 1 - \sum_k D_{ik} M_{ik}$. In view of the definition of M_{ik} , σ_i represents **the sum of the labels in i -th instance that are masked**, and the recovery bound of i -th instance ϵ_i becomes

$$\epsilon_i \leq \frac{\sigma_i^2}{(\sum_k D_{ik}^g M_{ik})^2}. \quad (10)$$

From Eq. (10), we can observe that its denominator is strictly greater than zero but less than one. And the smaller the σ_i is, the tighter the recovery bound ϵ_i is. When σ_i tends to zero, ϵ_i also tends to zero. This degrades the hidden label LDL problem to the learnable traditional LDL problem, which is an intuitive conclusion. *Combining all the proofs together, we conclude that LDL with hidden labels through our method is theoretically feasible.*

3.4 Optimizing using ADMM

Considering that the optimization objective contains complex constraints and non-smooth terms, we use ADMM (Alternating Direction Method of Multipliers) (Boyd et al. 2011) to solve Eq. (7). By breaking Eq. (7) into smaller pieces, each of which will be easier to handle. Specifically, Eq. (7) can be converted into the following equivalent forms by replacing the matrix \mathbf{D} with \mathbf{A} and \mathbf{B} :

$$\min_{\mathbf{A}, \mathbf{B}, \mathbf{D}} \frac{1}{2} \text{tr}(\mathbf{D}^\top \mathbf{G} \mathbf{D}) + \alpha \|\mathbf{A}\|_* \quad (11)$$

s.t. $\mathbf{D} \times \mathbf{1}_m = \mathbf{1}_n, \mathbf{D} \geq \mathbf{0}_{n \times m}, \mathbf{B} \in \text{Cons},$
 $\mathbf{D} - \mathbf{A} = \mathbf{0}, \mathbf{D} - \mathbf{B} = \mathbf{0}.$

Then, the solution of Eq. (11) can be accomplished by solving the following augmented Lagrange multiplier equation:

$$\min_{\mathbf{A}, \mathbf{B}, \mathbf{D}} \frac{1}{2} \text{tr}(\mathbf{D}^\top \mathbf{G} \mathbf{D}) + \alpha \|\mathbf{A}\|_* \quad (12)$$

+ $\langle \mathbf{\Lambda}, \mathbf{D} - \mathbf{A} \rangle + \frac{\rho}{2} \|\mathbf{D} - \mathbf{A}\|_F^2$
+ $\langle \mathbf{\Lambda}', \mathbf{D} - \mathbf{B} \rangle + \frac{\rho}{2} \|\mathbf{D} - \mathbf{B}\|_F^2$
s.t. $\mathbf{D} \times \mathbf{1}_m = \mathbf{1}_n, \mathbf{D} \geq \mathbf{0}_{n \times m}, \mathbf{B} \in \text{Cons},$

where $\mathbf{\Lambda} \in R^{n \times m}$ and $\mathbf{\Lambda}' \in R^{n \times m}$ are the Lagrange multipliers, $\rho \geq 0$ is a penalty parameter, and $\langle \cdot, \cdot \rangle$ returns the inner product of two matrices. Although ρ can increment in each loop, in our paper, we fix ρ at 2, which is sufficient for the objective to converge. To solve Eq. (12), the ADMM iteratively solves the following subproblems.

1) \mathbf{D} Subproblem: Removing unrelated terms regarding \mathbf{D} , \mathbf{D} subproblem is written as:

$$\min_{\mathbf{D}} \frac{1}{2} \text{tr}(\mathbf{D}^\top \mathbf{G} \mathbf{D}) + \langle \mathbf{\Lambda} \cdot \mathbf{D} - \mathbf{A} \rangle + \frac{\rho}{2} \|\mathbf{D} - \mathbf{A}\|_F^2 \quad (13)$$

+ $\langle \mathbf{\Lambda}' \cdot \mathbf{D} - \mathbf{B} \rangle + \frac{\rho}{2} \|\mathbf{D} - \mathbf{B}\|_F^2$
s.t. $\mathbf{D} \times \mathbf{1}_m = \mathbf{1}_n, \mathbf{D} \geq \mathbf{0}_{n \times m}.$

We use projected gradient descent (PGD) to solve Eq. (13). The gradient of the objective with respect to \mathbf{D} is:

$$\nabla_{\mathbf{D}} = \mathbf{G} \mathbf{D} + \mathbf{\Lambda} + \mathbf{\Lambda}' + \rho(\mathbf{D} - \mathbf{A}) + \rho(\mathbf{D} - \mathbf{B}). \quad (14)$$

We use Stochastic Gradient Descent (SGD) to update \mathbf{D} one time in a step and project it to probability simplex.

Although for every instance \mathbf{d}_i , in (Li and Chen 2024), the projection onto a probability simplex problem is solved

using an $O(m \log m)$ (Wang and Carreira-Perpinán 2013) algorithm by brute force searching through $[m]$ from the largest entry to the smallest one for a particular j satisfying the KKT condition, this projection method is not efficient. We use a much simpler projection function in $O(m)$:

Step I: Setting all the D_{ij} to 0 if $D_{ij} < 0$,

$$\text{Step II: Normalize } \mathbf{D} \text{ as } D_{ij} = D_{ij} / \sum_{k=1}^m D_{ik}. \quad (15)$$

We set a uniform distribution if $\sum_{k=1}^m D_{ik} = 0$. This projection method can be realized in $O(m)$, which is much faster than (Wang and Carreira-Perpinán 2013). At the same time, it can maintain proportional information and ultimately achieve better results in experiments.

2) **A Subproblem:** Removing irrelated terms regarding \mathbf{A} , \mathbf{A} subproblem can be rewritten into

$$\min_{\mathbf{A}} \frac{1}{2} \|\mathbf{A} - (\mathbf{D} + \frac{\mathbf{\Lambda}}{\rho})\|_F^2 + \frac{\alpha}{\rho} \|\mathbf{A}\|_*, \quad (16)$$

which has a closed-form solution by using a singular value thresholding operator (Cai, Candès, and Shen 2010).

3) **B Subproblem:** Removing irrelated terms regarding \mathbf{B} , \mathbf{B} subproblem can be solved by optimizing each row of \mathbf{B} ,

$$\begin{aligned} \min_{\mathbf{b}_i} \langle \lambda'_i, \mathbf{d}_i - \mathbf{b}_i \rangle + \frac{\rho}{2} \|\mathbf{d}_i - \mathbf{b}_i\|_F^2 \\ \text{s.t. } \mathbf{b}_i \in \{ \mathbf{e}_i \mid \mathbf{E} \in \text{Cons} \}, \end{aligned} \quad (17)$$

where the i -th row of \mathbf{B} is denoted as \mathbf{b}_i , so is $\lambda'_i, \mathbf{d}_i, \mathbf{m}_i$, and \mathbf{e}_i . Note that in Eq. (17), if $M_{ij} = 0$, no matter what k_i is, $k_i \cdot D_{ij}^o = B_{ij} \cdot M_{ij} = 0$. As a result, for those hidden positions, constraint is automatically satisfied. This problem turns out to be an unconstrained quadratic optimization, and

B_{ij} is directly set as $D_{ij} + \frac{\Lambda'_{ij}}{\rho}$. When it comes to observed positions, $B_{ij} = k_i \cdot D_{ij}^o$, where D_{ij}^o is a constant. Replacing all the B_{ij} with k_i , the optimization object is a quadratic equation about k_i , which is easy to deal with. Therefore, \mathbf{B} can be updated as follows:

$$B_{ij} = \begin{cases} D_{ij} + \frac{\Lambda'_{ij}}{\rho} & M_{ij} = 0, \\ \frac{\sum_{k=1}^m (\rho D_{ik} + \Lambda'_{ik}) \cdot D_{ik}^o \cdot M_{ik}}{\rho \sum_{k=1}^m (D_{ik}^o)^2 \cdot M_{ik}} D_{ij}^o & M_{ij} = 1. \end{cases} \quad (18)$$

Note that at least one entry in \mathbf{d}_i^o is not 0, so $\rho \sum_{k=1}^m (D_{ik}^o)^2 \cdot M_{ik} > 0$. As a result, Eq. (18) is always meaningful and B_{ij} can be updated correctly.

Furthermore, we update the Lagrange multiplier matrices as follows:

$$\begin{cases} \mathbf{\Lambda} = \mathbf{\Lambda} + \rho(\mathbf{D} - \mathbf{A}), \\ \mathbf{\Lambda}' = \mathbf{\Lambda}' + \rho(\mathbf{D} - \mathbf{B}). \end{cases} \quad (19)$$

See the overall pseudocode of the solution to Eq. (7) in Algorithm 1. Besides, the stopping criterion of the algorithm is that the number of iteration reaches 100, or the maximum of all the residuals of the optimized variables is less than

Algorithm 1: Optimization to Eq. (7)

```

1: Input:  $\mathbf{D}_o$ , the Laplacian matrix  $\mathbf{G}$ , mask  $\mathbf{M}$ 
2: Initialization:  $\mathbf{A} = \mathbf{B} = \mathbf{\Lambda}_1 = \mathbf{\Lambda}_2 = \mathbf{1}_{n \times m}$ ,  $\mathbf{D} = \mathbf{D}_o$ ,  $t = 1, \rho = 2$ 
3: while stopping criterion is not satisfied do
4:   Update  $\mathbf{D}$  by Eq. (14);
5:   Project  $\mathbf{D}$  by Eq. (15);
6:   Solve  $\mathbf{A}$  by Eq. (16);
7:   Update  $\mathbf{B}$  by Eq. (18);
8:   Update  $\mathbf{\Lambda}, \mathbf{\Lambda}'$  by Eq. (19);
9:    $t = t + 1$ ;
10: end while
11: return  $\mathbf{D}$ 

```

10^{-3} , i.e., $\max(\|\mathbf{D} - \mathbf{A}\|_\infty, \|\mathbf{D} - \mathbf{B}\|_\infty) < 10^{-3}$, where $\|\cdot\|_\infty$ denotes the infinity norm of the matrix.

4 Experiment

4.1 Datasets

In this paper, we evaluate our method on 12 real-world datasets covering fields of biology, facial expression, natural scene, emotion recognition and movie. The details of the statistics of these datasets can be found in (Geng 2016; Peng et al. 2015; Li and Deng 2019).

To simulate the presence of hidden labels in the datasets, we first make these datasets randomly missing (Xu and Zhou 2017) and set the missing d_x^y to 0. We vary the missing rate ω from 40% to 80%. Then, we check the randomly missing datasets and ensure that for each instance \mathbf{x} , at least one label description has a value greater than 0 (i.e., $\exists j$ such that $d_x^{y_j} > 0$). In HidLDL, the input is a seemingly complete label distribution dataset, so in the second step, we normalize missing datasets into the final hidden setting using Eq. (3).

4.2 Settings and Compared Methods

To validate the effectiveness of our method, we design two experimental configurations. In the first setting, we use a hidden dataset and features of instances together to restore the true label distribution. Difference between the ground-truth and the recovered distribution matrix will be measured. In the second setting, to further verify the quality of the recovered label distribution, we train each recovered label distribution using the same simple LDL algorithm, such as SA-BFGS (Geng 2016). We then use each trained model for prediction. The training and testing sets are partitioned with ratios of 0.8 and 0.2, respectively. We repeat each experiment 5 times and report the average results.

We compare our method with six methods. Two of them are state-of-the-art methods designed for IncomLDL including IncomLDL-admm (abbreviated as InLDL-a) (Xu and Zhou 2017) and WInLDL (Li and Chen 2024), which can still be used for the hidden setting. Two state-of-the-art LDL methods named LDL-LRR (Jia et al. 2021) and LDL-DPA (Jia et al. 2023) are included. We also include two baselines named SA-IIS and PT-Bayes (Geng 2016). For comparison methods, we use default parameters suggested in their original papers, with the exception that we adjust the regulariza-

Dataset	Ours	InLDL-a	WInLDL	SA-IIS	LDL-LRR	PT-Bayes	LDL-DPA
alpha	0.4718 ± .0011	0.6878 ± .0010 [•]	0.9660 ± .0179 [•]	0.9466 ± .0084 [•]	0.7078 ± .0056 [•]	1.9636 ± .0673 [•]	1.0037 ± .0117 [•]
cdc	0.3803 ± .0044	0.6493 ± .0023 [•]	0.8870 ± .0172 [•]	0.8590 ± .0154 [•]	0.6607 ± .0050 [•]	1.6762 ± .0708 [•]	0.9007 ± .0038 [•]
cold	0.1746 ± .0005	<u>0.2442 ± .0013</u> [•]	0.2800 ± .0049 [•]	0.2918 ± .0045 [•]	0.2519 ± .0045 [•]	0.4631 ± .0551 [•]	0.3068 ± .0086 [•]
dtc	0.1237 ± .0008	<u>0.1734 ± .0008</u> [•]	0.2200 ± .0058 [•]	0.2328 ± .0063 [•]	0.1825 ± .0062 [•]	0.4363 ± .0410 [•]	0.2499 ± .0032 [•]
elu	0.3458 ± .0072	<u>0.5899 ± .0011</u> [•]	0.8038 ± .0166 [•]	0.7791 ± .0151 [•]	0.6029 ± .0073 [•]	1.4970 ± .0896 [•]	0.8156 ± .0223 [•]
spo	0.3250 ± .0064	<u>0.5162 ± .0011</u> [•]	0.5619 ± .0042 [•]	0.5646 ± .0075 [•]	0.5202 ± .0020 [•]	0.8059 ± .0204 [•]	0.5886 ± .0072 [•]
SJAFFE	0.5983 ± .0235	0.8592 ± .0535 [•]	2.0855 ± .1299 [•]	<u>0.8421 ± .0339</u> [•]	0.8891 ± .0050 [•]	0.9838 ± .0659 [•]	2.0990 ± .1017 [•]
Scene	6.4190 ± .0026	6.7305 ± .0113 [•]	6.7673 ± .0186 [•]	6.7141 ± .0078 [•]	6.7761 ± .0050 [•]	7.2195 ± .1133 [•]	6.7538 ± .0139 [•]
Movie	0.7494 ± .0041	<u>1.1138 ± .0047</u> [•]	1.5030 ± .0083 [•]	1.6548 ± .0117 [•]	1.1415 ± .0068 [•]	4.2956 ± .1179 [•]	1.4271 ± .0389 [•]
SBU	0.7028 ± .0086	0.8375 ± .0048 [•]	0.8278 ± .0103 [•]	0.8686 ± .0041 [•]	0.8998 ± .0036 [•]	0.8737 ± .0088 [•]	0.8552 ± .0062 [•]
Emo	2.4372 ± .1010	3.7105 ± .0233 [•]	3.7991 ± .0087 [•]	3.7944 ± .0125 [•]	3.7732 ± .0046 [•]	3.8714 ± .0415 [•]	3.7921 ± .0150 [•]
RAF	3.0345 ± .0095	4.6553 ± .0079 [•]	5.2983 ± .0065 [•]	2.9245 ± .0096 [◦]	2.9998 ± .0054 [◦]	3.1821 ± .0440 [•]	2.9374 ± .0112 [•]

Table 1: Canberra (the lower the better) results for the **recovery** setting on all datasets when missing rate $\omega = 50\%$. The value is shown in mean \pm std form. Bold and underlined indicate the best and second best results, respectively. Under single-tailed paired t-test at a significance level of 0.05, \bullet means our method’s performance is statistically significantly better than the compared method and \circ means the compared method’s performance is statistically significantly better than ours.




Figure 2: The visualization of two typical recovery results on the Movie (left) and spo (right) dataset.

tion parameter for IncomLDL-admm from $2^{\{-10,-9,\dots,9,10\}}$. For our method, the regularization parameter α is selected from $2^{\{-10,-9,\dots,9,10\}}$ based on the recover performance, while σ and ρ are set to be 1 and 2, respectively. In K-nearest neighbors algorithm, the number of neighbors is set as the number of labels. Each method will face a matrix with exactly the same missing entries in each experiment.

We use five metrics for the LDL with Hidden Labels problem, including Chebyshev, Clark, Canberra, Cosine and Intersection. Details of them can be found in (Geng 2016). Among them, Canberra, Clark and Canberra measure the distance between two vectors, thus they are the lower the better. Cosine and Intersection measure the similarity between two vectors, thus they are the higher the better.

4.3 Recover Results and Discussions

In the following, we report the recover results of different methods at different missing rate. Due to space limitation, here we only list representative results. The Canberra (the lower the better) results on all the dataset with $\omega = 50\%$ are shown in Table 1.

To conduct significance analysis, we conduct the paired t-test, which is a useful statistical test for comparing two methods (Box 1987). For each dataset, we perform 6 separate single-tailed paired t-tests, comparing our method with 6 compared methods individually. We mark the compared method with a black dot \bullet in Table 1 if our method’s performance is statistically significantly better than the compared

Metric	Method	cold	dtc	spo	SBU	Emo	RAF
Clark \downarrow	Ours	0.099	0.072	0.163	0.327	1.249	1.486
	w/o Cons	0.147	0.102	0.256	0.413	1.679	1.611
	w/o TN	0.123	0.099	0.189	0.331	1.538	1.583
Cosine \uparrow	Ours	0.993	0.997	0.989	0.951	0.825	0.783
	w/o Cons	0.987	0.993	0.975	0.918	0.657	0.642
	w/o TN	0.989	0.996	0.985	0.947	0.685	0.664

Table 2: Ablation Results on 6 Datasets. \uparrow (\downarrow) indicates the higher (lower) the better.

method at a significance level of 0.05.

By analyzing the experimental results, we can arrive at the following three conclusions:

- Our method achieves the lowest average rank in terms of all five metrics. As shown in Table 1, we win 69 times out of 72 Canberra comparisons, with a **95.8%** rate to win, demonstrating that our method is superior to other compared methods. More importantly, our advantages are statistically significant on these 69 cases according to t-test. For detailed metrics, our improvement is particularly evident in Chebyshev, Clark and Canberra.
- IncomLDL methods are not much better than LDL methods in the recovery setting. WInLDL, although reported to be faster and more effective in the incomplete setting (Li and Chen 2024), performs quite poorly. InLDL-a, although performs relatively stably, is actually close to LDL-LRR in metrics’ values.
- As the missing rate increases, the performance metrics of almost all methods become worse. Our method has an outstanding advantage when the missing rate is between 40% and 60%.

Fig. 2 shows typical recovery results on two datasets. We selected two methods, InLDL-a and LDL-LRR, which are reported to have relatively better performance, to compare with our method. According to the visualization, it can be seen that our method (**red line**) is closer to the ground truth (**black line**). Meanwhile, our trend has greater consistency with the ground truth, which is helpful to identify labels with the largest and smallest description degrees.

Dataset	Ours	InLDL-a	WInLDL	SA-IIS	LDL-LRR	PT-Bayes	LDL-DPA
alpha	0.6890 ± .0012	0.6936 ± .0022 [•]	1.0180 ± .0094 [•]	0.9681 ± .0138 [•]	0.7073 ± .0101 [•]	1.8878 ± .0996 [•]	1.0802 ± .0173 [•]
cdc	0.6464 ± .0020	<u>0.6506 ± .0056</u>	0.9562 ± .0250 [•]	0.8731 ± .0175 [•]	0.6641 ± .0095 [•]	1.7754 ± .1194 [•]	0.9767 ± .0178 [•]
cold	0.2329 ± .0005	0.2369 ± .0021 [•]	0.2930 ± .0117 [•]	0.2834 ± .0071 [•]	0.2466 ± .0038 [•]	0.5084 ± .0815 [•]	0.3148 ± .0128 [•]
dtc	0.1635 ± .0002	<u>0.1674 ± .0010</u>	0.2370 ± .0120 [•]	0.2351 ± .0149 [•]	0.1916 ± .0126 [•]	0.4543 ± .0412 [•]	0.2714 ± .0215 [•]
elu	0.5701 ± .0009	<u>0.5771 ± .0035</u>	0.8423 ± .0178 [•]	0.7835 ± .0119 [•]	0.5931 ± .0057 [•]	1.5758 ± .1034 [•]	0.8857 ± .0186 [•]
spo	0.5169 ± .0011	<u>0.5197 ± .0039</u>	0.5956 ± .0137 [•]	0.5702 ± .0132 [•]	0.5243 ± .0048 [•]	0.8229 ± .0539 [•]	0.6182 ± .0193 [•]
SJAFFE	0.7935 ± .0182	0.8940 ± .0263 [•]	2.0713 ± .0657 [•]	0.8980 ± .0248 [•]	0.9120 ± .0093 [•]	1.0189 ± .0281 [•]	1.9060 ± .1365 [•]
Scene	6.8070 ± .0137	6.8531 ± .0254 [•]	6.9508 ± .0242 [•]	<u>6.8116 ± .0127</u>	6.8298 ± .0083 [•]	7.2288 ± .0596 [•]	6.8700 ± .0238 [•]
Movie	1.0787 ± .0080	1.2025 ± .0089 [•]	1.7374 ± .0250 [•]	1.9103 ± .0234 [•]	<u>1.1610 ± .0082</u>	4.2877 ± .0859 [•]	1.4859 ± .0210 [•]
SBU	0.8411 ± .0076	0.8716 ± .0053 [•]	0.8328 ± .0092 [◦]	0.9017 ± .0037 [•]	0.9142 ± .0007 [•]	0.8988 ± .0071 [•]	0.8892 ± .0094 [•]
Emo	3.7057 ± .0058	3.9040 ± .0219 [•]	3.8321 ± .0096 [•]	4.2095 ± .0456 [•]	3.8376 ± .0373 [•]	3.8614 ± .0815 [•]	3.9123 ± .0411 [•]
RAF	3.0544 ± .0043	2.8871 ± .0133 [◦]	4.3674 ± .0360 [•]	3.0108 ± .0123 [◦]	3.0033 ± .0108 [◦]	3.0488 ± .0245	2.9869 ± .0151 [◦]

Table 3: Canberra (the lower the better) results for the **predictive** setting on all datasets when missing rate $\omega = 50\%$. The value is shown in mean \pm std form. Bold and underlined indicate the best and second best results, respectively. Under single-tailed paired t-test at a significance level of 0.05, \bullet means our method’s performance is statistically significantly better than the compared method and \circ means the compared method’s performance is statistically significantly better than ours.

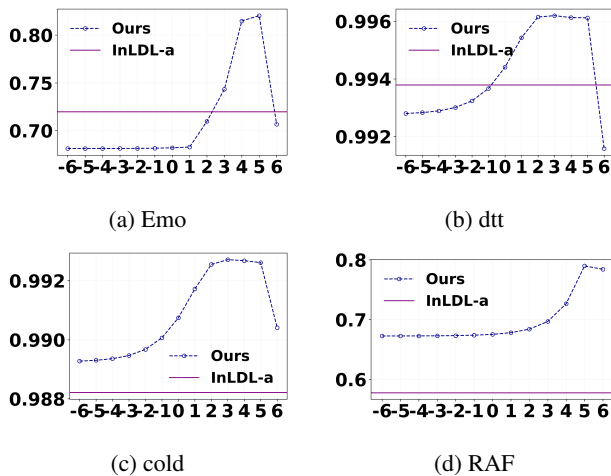


Figure 3: Comparison of our method and IncomLDL-admm on Cosine (the higher the better) and the horizontal axis represents $\log_2(\alpha)$ with the missing rate $\omega=50\%$.

Ablation Studies We perform ablation experiments on six datasets to validate the importance of our proportional constraint and global dependency. As shown in Table 2, “w/o Cons” method means our method without proportional constraint, while “w/o TN” means our method without the trace norm (TN) term (i.e., $\alpha = 0$). At 50% missing rate, “Ours” has a huge improvement in all the evaluation metrics compared to the “w/o Cons”, indicating that it is critical to maintain the proportional relationship between labels. Meanwhile, compared to “w/o TN”, there has also been an improvement.

Sensitivity Analysis Our method involves one parameter: α in the trace norm regularization. The analysis of the parameter α is shown in Fig. 3. To demonstrate the performance of our method, we set the IncomLDL-admm method to achieve its best tuned results. From Fig. 3, it can be seen that when $\alpha < 2^3$, our method shows a stable growth in performance as α increases. When $\alpha > 2^5$, the performance of our method deteriorates rapidly. As a result, when α is set

between 2^3 and 2^5 , relatively good results can be achieved.

4.4 Predict Results and Discussions

We present a representative result in Table 3 with a 50% missing rate for the metric Canberra, while the results for other missing rates and metrics are similar. Similar to the recovery setting, the single-tailed paired t-test is conducted to analyze the significant advantages of our method.

From the reported results, we observe that:

- Our method achieves the lowest average rank in terms of all five metrics. Taking Table 3 as an example, out of 12 datasets, we rank 1st 10 times and 2nd once. Out of 72 comparisons, we win 66 times. With a winning rate of over **91.6%**, our method performs better than most compared methods. And it is worth noting that our method has a significant advantage most of the time.
- Methods focusing on the relationship between labels (such as LDL-LRR and InLDL-a) can capture advantageous information in prediction even if their recovery results are poor.
- Our method is less affected by the missing rate in the predictive scenarios. As the missing rate increases, our method is robust and deteriorates very slowly.

5 Conclusion

This paper gives a creative answer to the question: “What does real-world incomplete data in IncomLDL actually look like and how to deal with it”. Specifically, we clarify the irrationality of the training dataset in IncomLDL, and repair it by normalization. We then proposed a new method to solve this brand-new HidLDL problem, which emphasizes maintaining the observed labels’ proportion during the recovery process. We then theoretically prove the feasibility of our method in learning from hidden labels. Furthermore, we theoretically provide the recovery bound of our method, demonstrating that it is feasible to learn from hidden labels. Extensive experiments validate advantages of our method against other IncomLDL and LDL methods in HidLDL.

Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grants U24A20322, 62576094 and 62422118. This work is also supported by Hong Kong UGC under grants UGC/FDS11/E03/24, UGC/FDS11/E03/25, and Hong Kong Research Grants Council under Grant 11219324. This research work is also supported by the Big Data Computing Center of Southeast University.

References

- Box, J. F. 1987. Guinness, Gosset, Fisher, and small samples. *Statistical science*, 45–52.
- Boyd, S.; Parikh, N.; Chu, E.; Peleato, B.; Eckstein, J.; et al. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1): 1–122.
- Cai, J.-F.; Candès, E. J.; and Shen, Z. 2010. A singular value thresholding algorithm for matrix completion. *SIAM Journal on optimization*, 20(4): 1956–1982.
- Chen, S.; Wang, J.; Chen, Y.; Shi, Z.; Geng, X.; and Rui, Y. 2020. Label Distribution Learning on Auxiliary Label Space Graphs for Facial Expression Recognition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, 13981–13990. Computer Vision Foundation / IEEE.
- Gao, B.; Zhou, H.; Wu, J.; and Geng, X. 2018. Age Estimation Using Expectation of Label Distribution Learning. In Lang, J., ed., *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 712–718. ijcai.org.
- Geng, X. 2016. Label Distribution Learning. *IEEE Transactions on Knowledge and Data Engineering*, 28(7): 1734–1748.
- Geng, X.; and Hou, P. 2015. Pre-release prediction of crowd opinion on movies by label distribution learning. In *IJCAI*, 3511–3517. Citeseer.
- Geng, X.; Yin, C.; and Zhou, Z.-H. 2013. Facial age estimation by learning from label distributions. *IEEE transactions on pattern analysis and machine intelligence*, 35(10): 2401–2412.
- Jia, X.; Li, W.; Liu, J.; and Zhang, Y. 2018. Label distribution learning by exploiting label correlations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Jia, X.; Li, Z.; Zheng, X.; Li, W.; and Huang, S.-J. 2019a. Label distribution learning with label correlations on local samples. *IEEE Transactions on Knowledge and Data Engineering*, 33(4): 1619–1631.
- Jia, X.; Qin, T.; Lu, Y.; and Li, W. 2023. Adaptive weighted ranking-oriented label distribution learning. *IEEE Transactions on Neural Networks and Learning Systems*.
- Jia, X.; Shen, X.; Li, W.; Lu, Y.; and Zhu, J. 2021. Label distribution learning by maintaining label ranking relation. *IEEE Transactions on Knowledge and Data Engineering*, 35(2): 1695–1707.
- Jia, X.; Zheng, X.; Li, W.; Zhang, C.; and Li, Z. 2019b. Facial emotion distribution learning by exploiting low-rank label correlations locally. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 9841–9850.
- Jia, Y.; Liu, H.; Hou, J.; Kwong, S.; and Zhang, Q. 2022. Semisupervised Affinity Matrix Learning via Dual-Channel Information Recovery. *IEEE Trans. Cybern.*, 52(8): 7919–7930.
- Jia, Y.; Tang, J.; and Jiang, J. 2024. Label Distribution Learning from Logical Label. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*, 4228–4236.
- Jia, Y.; Tao, S.; Wang, R.; and Wang, Y. 2024. Ensemble Clustering via Co-Association Matrix Self-Enhancement. *IEEE Trans. Neural Networks Learn. Syst.*, 35(8): 11168–11179.
- Kou, Z.; Qin, S.; Wang, H.; Wang, J.; Xie, M.; Chen, S.; Jia, Y.; Liu, T.; Sugiyama, M.; and Geng, X. 2025a. Label Distribution Learning with Biased Annotations Assisted by Multi-Label Learning. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2025, Montreal, Canada, August 16-22, 2025*, 5545–5553.
- Kou, Z.; Wang, J.; Jia, Y.; Liu, B.; and Geng, X. 2025b. Instance-Dependent Inaccurate Label Distribution Learning. *IEEE Trans. Neural Networks Learn. Syst.*, 36(1): 1425–1437.
- Li, S.; and Deng, W. 2019. Blended emotion in-the-wild: Multi-label facial expression recognition using crowd-sourced annotations and deep locality feature learning. *International Journal of Computer Vision*, 127(6): 884–906.
- Li, X.; and Chen, S. 2024. No regularization is needed: efficient and effective incomplete label distribution learning. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 4470–4478.
- Li, Y.; Li, Z.; and Jia, Y. 2025. Complementary Label Learning with Positive Label Guessing and Negative Label Enhancement. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Liang, L.; Lin, L.; Jin, L.; Xie, D.; and Li, M. 2018. SCUT-FBP5500: A diverse benchmark dataset for multi-paradigm facial beauty prediction. In *2018 24th International conference on pattern recognition (ICPR)*, 1598–1603. IEEE.
- Liu, W.; Wang, H.; Shen, X.; and Tsang, I. W. 2021. The emerging trends of multi-label learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(11): 7955–7974.
- Ma, H.; Lu, N.; Mei, J.; Guan, T.; Zhang, Y.; and Geng, X. 2023. Label distribution learning for scene text detection. *Frontiers Comput. Sci.*, 17(6): 176339.
- Peng, K.-C.; Chen, T.; Sadovnik, A.; and Gallagher, A. C. 2015. A mixed bag of emotions: Model, predict, and transfer emotion distributions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 860–868.

- Ren, Y.; and Geng, X. 2017. Sense Beauty by Label Distribution Learning. In Sierra, C., ed., *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, 2648–2654. ijcai.org.
- Tang, J.; and Jia, Y. 2025. Concentration Distribution Learning from Label Distributions. In *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, Canada, July 13-19, 2025*.
- Tsoumakas, G.; and Katakis, I. 2008. Multi-label classification. *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications: Concepts, Methodologies, Tools, and Applications*, 3: 64.
- Wang, J.; and Geng, X. 2021. Label distribution learning by exploiting label distribution manifold. *IEEE transactions on neural networks and learning systems*, 34(2): 839–852.
- Wang, J.; Zhang, F.; Jia, X.; Wang, X.; Zhang, H.; Ying, S.; Wang, Q.; Shi, J.; and Shen, D. 2022. Multi-Class ASD Classification via Label Distribution Learning with Class-Shared and Class-Specific Decomposition. *Medical Image Analysis*, 75: 102294.
- Wang, W.; and Carreira-Perpinán, M. A. 2013. Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application. *arXiv preprint arXiv:1309.1541*.
- Wen, C.; Zhang, X.; Yao, X.; and Yang, J. 2023. Ordinal Label Distribution Learning. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 23424–23434.
- Xie, K. Y.; Wang, J.; Jia, Y.; Shi, B.; and Geng, X. 2026. RankMatch: A Novel Approach to Semi-Supervised Label Distribution Learning Leveraging Inter-label Correlations. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Xu, M.; and Zhou, Z.-H. 2017. Incomplete Label Distribution Learning. In *IJCAI*, 3175–3181.
- Xu, S.; Shang, L.; Shen, F.; Yang, X.; and Pedrycz, W. 2025. Incomplete label distribution learning via label correlation decomposition. *Information Fusion*, 113: 102600.
- Yang, X.; Gao, B.-B.; Xing, C.; Huo, Z.-W.; Wei, X.-S.; Zhou, Y.; Wu, J.; and Geng, X. 2015. Deep label distribution learning for apparent age estimation. In *Proceedings of the IEEE international conference on computer vision workshops*, 102–108.
- Zhang, M.; and Zhou, Z. 2014. A Review on Multi-Label Learning Algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 26(8): 1819–1837.
- Zhu, X. 2005. *Semi-supervised learning with graphs*. Carnegie Mellon University.